

A METHOD OF COMMUNICATION IN A TIME DIVISION DUPLEX (TDD)
SATELLITE COMMUNICATION SYSTEM

This invention relates to a method of communication in a time division duplex
5 (TDD) satellite communication system.

Time Division Duplex simplifies the hardware of satellites by eliminating the
complex, heavy and costly diplexer filters required to facilitate the use of frequency
division duplex (FDD). This, in turn, allows more complex RF structures to be
implemented, including the use of larger numbers of elements in phased array antenna
10 structures.

However, traditionally, the use of TDD has been associated with short to
medium range terrestrial links where the propagation delays can be kept small in
comparison with the length of the frames for TDD operation. This has been a
requirement because conventional wisdom was that the TDD frames needed to
15 incorporate guard periods equal to the maximum round trip propagation delay in order
to avoid interference between uplink and downlink under worst case conditions.
The above constraint would clearly be prohibitive in a satellite application given a
minimum round trip delay of order 240 ms. Thus, it is necessary to provide
improvements for TDD in a satellite context to allow operation with high efficiency
20 (minimal guard time) and a low marginal delay (short TDD frame).

In accordance with the present invention, a method of communication in a time
division duplex (TDD) satellite communication system comprising at least one satellite
and a plurality of terrestrial terminals comprises allocating time division multiple
access (TDMA) time slots for transmission between the satellite and any one of the
25 plurality of terminals, such that for any given terminal, transmit time slots for
transmission to the satellite and receive time slots for reception from the satellite are
separated in time; wherein an assigned time delay between transmit and receive time
slots at the any one terminal is small compared with a round trip propagation delay; and
wherein, if the transmit time slot for one terminal causes a transmission from that one
30 terminal to be received at another terminal overlapped in time with a receive time slot
allocated for the other terminal, then those two terminals are spaced apart in distance,
sufficiently, such that interference between the two terminals is minimised.

Preferably, signals between the terminals and the satellite are synchronised at the satellite.

Preferably, alternate time slots at the satellite are used for transmission and reception.

5 Preferably, the terminals use navigational information to estimate their propagation delay to the satellite; and thus to determine the time required to transmit into an allocated time slot.

Preferably, the satellite transmits ephemeris data to the terminals to aid in determining the propagation delay.

10 Preferably, the position of each terminal is determined by the satellite, using location data provided by each terminal.

Preferably, downlink timeslots are allocated to terminals at random.

Preferably, uplink timeslots are allocated in order to avoid a transmission at one terminal being received by another terminal at a time for which the other terminal has
15 been allocated a receive time slot.

Preferably, terminal receive time slots are allocated randomly; wherein allocation of terminal transmit time slots includes the steps of: calculating the minimum distance between a transmitting terminal and a receiving terminal which receives the transmission; repeating this calculation for all terminal transmit time slots;
20 repeating the calculation for all terminals; calculating the resulting interference if each terminal used its worst terminal transmit time slot; ranking the terminals according to which cause the worst interference with another terminal; and starting from the worst terminal, allocating the best time slot for that terminal, discarding terminal transmit time slots where transmit and receive time slots overlap in the same terminal.

25 An example of a method of communication in a TDD satellite communication system comprising at least one satellite and a plurality of terrestrial terminals will now be described with reference to the accompanying drawings in which:-

Figure 1 illustrates ideal timing for a basic TDD system;

Figure 2 illustrates worst case timing for a basic TDD system;

30 Figure 3 illustrates distributed timing for a basic TDD system;

Figure 4 illustrates ideal timing for a TDD/TDMA method according to the present invention;

Figure 5 illustrates worst case timing for a TDD/TDMA method according to the present invention;

Figure 6 illustrates worst case timing for an offset TDD method applied to the TDD/TDMA method according to the present invention;

5 Figure 7 illustrates the best case using a reversed TDD/TDMA according to the present invention;

Figure 8 illustrates the worst case for a reversed TDD/TDMA according to the present invention;

10 Figure 9 illustrates the occurrence of interfering terminals and the distance between them for a non optimised interference relationship for the maximum coverage of the satellite;

Figure 10 illustrates the occurrence of interfering terminals and the distance between them for an optimised interference relationship for the maximum coverage of the satellite;

15 Figure 11 illustrates the occurrence of interfering terminals and the distance between them for a non optimised interference relationship for a spot beam of 60km radius from the satellite;

20 Figure 12 illustrates the occurrence of interfering terminals and the distance between them for an optimised interference relationship for a spot beam of 60km radius from the satellite;

Figure 13 illustrates the cumulative distribution of minimum distance in km having run a simulation for the Fig.11 and Fig.12 examples for a series of different random terminal positions.

25 Figure 14 shows satellite distance geometries for inter-satellite interference.

In a simple TDD scheme, the satellite round trip propagation delay is large, so it is highly desirable that any physical layer related delays do not add to this significantly. For this reason the TDD frame length is chosen to be a fraction of the propagation delay. An example in which the one way propagation delay is exactly five times the
30 frame length is illustrated in Fig.1. Transmit time slots 1 from the satellite 2 are

interleaved with reception time slots 3 to receive signals from terrestrial terminals 4 with respective receive and transmit time slots 5, 6. In this case the situation is ideal as every satellite transmission coincides exactly with a terminal receive time slot and vice versa.

5 However, the propagation delay is not always an exact number of multiples of the frame length. Fig.2. illustrates a case in which the one way propagation delay is 4.75 times the frame length. In this case, the timing established by the satellite 2 forces the terminals 4 to receive and transmit at the same time. This is only possible if two
10 separate frequencies are used which is undesirable, as half of the time is wasted altogether and the other time is unusable. This problem can be alleviated by sharing the timing shifts between the satellite 2 and the terminal 4, as shown in Fig. 3 so that at least some of the time is usable. However, this leaves the capacity effectively reduced to 50%.

 By introducing TDMA, the situation is improved considerably. The case for
15 ideal timing with four slot TDMA is shown in Fig.4. Here the first slot 7, 8 in the satellite and terminal respectively is used for both uplink and downlink 9, 10. The situation is not dissimilar to that shown in Fig. 1, except that the transmission and reception times are shorter.

 Fig. 5 illustrates the situation with TDMA where the one way propagation delay
20 is 4.75 times the frame length. Here it can be seen, in the same way as for Fig.2, that the terminal needs to transmit and receive 11 at the same time. However, with TDMA there is additional flexibility to change the timing to avoid such conflicts. This is illustrated in Fig. 6 where the first slot 7, 8 is used for transmission and the third slot 12, 13 for reception. However, this approach, whilst solving the problem is relatively
25 inflexible.

 A more flexible solution to the problem of avoiding both transmit and receive being required together is to alter the switching of the TDD and TDMA by setting the innermost switching to be TDD switching and the outermost to be TDMA switching. In this case the structure is essentially the same as in Fig. 1 except that the source and
30 destinations from successive time slots are different. For simplicity, the examples given show four time slots for TDMA, but in practice there are more likely to be between 10 and 100 slots.

In Fig.7 the example shows a series of channels, labelled A, B, C & D, in both the satellite and the terminal. The downlink 10 time slot is followed by an uplink time slot 9 at the satellite for each successive channel A, B, C, D. For simplicity, the terminals 4 are assumed to be all at the same range to the satellite 2. Nevertheless, because of the large propagation delay, the uplink and downlink slots 5, 6 have become scrambled, but this is of no consequence.

In the case where the propagation delay is the worst case, there is a similar situation to Fig. 2, but with an important difference, as illustrated in Fig. 8. Here, as before, transmission and reception is required on the earth at the same time. However, these requirements are not at the same place. Thus, for example, terminal D must receive while terminal A is transmitting, or terminal A receives whilst terminal B is transmitting. If the range between these two terminals is large enough then the interference path will be negligible, so this is feasible. Moreover, the order of the uplink time slots can be scrambled so as to maximise the distance between mutually interfering terminals.

In practice, the terminals are spread over an area large enough such that the propagation delays to the satellite vary considerably, i.e. by many time slots. Far from being a problem, this effectively randomises the interference and availability of uplink time slots that do not overlap with the corresponding downlink time slot. The system then operates as follows. All timing is synchronous at the satellite, but the timings on the ground depend on the propagation delays. Thus, terrestrial terminals are set to receive when their time slot reaches their location, but they must transmit at a time necessary to ensure that their signal is received into the allocated uplink time slot at the satellite.

Switching between uplink and downlink happens following transmission or reception for individual users, so instead of transmitting all users' downlinks followed by all users' uplinks, the uplink and downlink time slots 1, 3 are interleaved. The slot time is very short, of the order of 100 μ s, compared with the round trip propagation delay. The up and downlink time slots are assigned independently to avoid terminals being required to receive and transmit at the same time and to maximise the range between interfering terminals on the ground. Thus, timing is synchronous at the satellite but appears arbitrary on the ground.

A simulation has been written to evaluate the potential for minimising interference and the available capacity. The results indicate that 100% capacity (i.e. no wasted time slots) is achievable with low interference on the ground. A method for providing control channels has also been determined. This takes up two channels. Thus
 5 for a system based on 100 channels the overhead is only 2%. With 100 μ s time slots the TDMA frame length would be 20 ms for 100 hundred channels. Thus the marginal impact on service delay (given 240 ms round trip satellite delay) is negligible. Interference between satellites is acceptable for three satellite global coverage. It is possible to use favourable synchronisation to avoid interference between adjacent
 10 satellites at shorter range. It is assumed that satellite navigation is used to provide the locations of the terminals. This is used to assist in the setting of timing advance (e.g. for using the random access channel). The locations are also relayed to the satellite so that it can use them to assist in setting up non interfering slot allocations.

The simulation of performance of the system was carried out using the
 15 following assumptions. A geostationary satellite was positioned above a point over the equator. A coverage area with circular perimeter was identified directly underneath the satellite. Initially this area was the maximum that gave satellite visibility (satellite elevation 0° at perimeter). A fixed number of terminals were deployed randomly with uniform distribution over the ground. The great circle distances between all terminal
 20 pairs were computed. Downlink time slots were assigned arbitrarily to each terminal and uplink time slots were assigned, initially in the same order as the downlink time slots, but with provision for some slots to be shifted if they require a terminal to have overlapped receive and transmit times.

The distances between terminals were used to determine any cases where one
 25 terminal's transmission, delayed by the propagation time, overlapped with another terminal's reception time. For those cases, the distances between terminals were noted. A simulation was run for 100 terminals, with a slot length of 100 μ s. Fig. 9 is a surface plot giving an indication of the result. The X and Y co-ordinates identify a pair of terminals. If the Z value is zero at a particular position then there is no interference
 30 between the terminals. Otherwise, the height indicates the distance between the interfering terminals. For the example shown, the worst case or minimum distance between interferers was 926 km. This is more than adequate for acceptable interference. However, this can be further improved upon.

In the example shown in Fig.9, the uplink time slots were allocated without reference to ground-to-ground interference effects. A more sophisticated assignment allocation has also been developed. This is based on the assumption that the satellite knows the positions of the terminals. This is not unreasonable if the terminals have GPS receivers and can relay their position information to the satellite on some control channel. The positions are only needed to a coarse degree so even with relatively high mobility this information would only need to be transferred infrequently. The principles of the optimising algorithm are as follows. First the downlink slots are assigned, which is assumed to be done on an arbitrary basis. Then, the effects for a particular terminal of assigning a time slot are examined in terms of the potential for that terminal interfering with other terminals' downlinks. The worst case interference that results is computed, i.e. the minimum distance between the transmitting terminal and the interfered terminal. This complete process is repeated for all possible time slots and for all terminals.

For every terminal, the worst case interference that will result if that terminal must use its worst time slot is determined and the terminals are ranked in order from the terminal that would have the worst interference if using its worst time slot, to the terminal that would have the least bad interference if using its worst time slot. Then, starting with the worst terminal, each terminal is assigned the best available time slot. In deciding on available time slots, any time slot resulting in the terminal having overlapped receive and transmit times is ignored.

An algorithm defining the optimising algorithm comprising the steps set out above was applied to the example of Fig. 9 (actual terminal positions) and the interference relationships derived, as shown in Fig. 10. This clearly shows that there are far fewer interferers. The minimum range to an interferer is now 4,670 km. This represents an improvement by a factor of over 5:1.

Given the ranges available without optimisation this improvement may seem academic. However, one of the advantages of using TDD is to facilitate greater use of phased array antennas and so allow spot beams. All time slots fall within a spot beam coverage area and this area is far smaller than the entire footprint of the geostationary satellite. For example, with a footprint of about 60 km, Fig. 11 shows the results for a typical scenario. The minimum distance between interferers is 6.5 km and occasionally much lower distances can arise. Fig. 12 illustrates the effects of applying optimisation.

The minimum distance is now 57.5 km, which is a significant fraction of the diameter of the cell.

To obtain more statistically significant results the simulation was run 1,000 times with different random positions for the terminals. The cumulative distributions of minimum range in km, both with and without the optimisation, were then plotted in Fig. 13. The optimised case is significantly better than the un-optimised case. There are a small number of cases where the range may be unacceptable. Further improvement of the algorithm is expected to deal with these cases. It should also be noted that the cell size is very small, so this represents a worse than worst case situation. A single satellite would need over 15,000 spot beams this size to cover 1/3 of the earth and even allowing for covering land masses only, this still represents a huge number.

One potential problem with TDD is that not only can terminals interfere with terminals, but satellites with satellites. For the case of 3 geostationary satellites covering the earth this is not a problem because the mutual ranges are greater than the satellite to ground ranges. Specifically, simple geometric considerations show that the worst case (i.e. for terminals on the edge of the coverage region) range ratio is about 1.85:1 corresponding to 5.4 dB in free space. Given the additional effects due to antenna patterns, inter satellite interference will not be a problem.

The above conclusion is satisfactory but unfortunate in that it appears to limit operation to cases where there are few satellites. In fact it is possible to bring the satellites closer together by ensuring that the propagation delays between adjacent satellites are whole multiples of the TDD period (i.e. twice the slot period – 200 μ s in our example corresponding to multiples of about 6 km). If this is achieved then transmissions from one satellite will arrive completely overlap with transmission slots for its neighbouring satellite. For non neighbouring satellites this relationship will gradually break down with increasing neighbour distance because the geometry of a circle will gradually take over, until $D \neq 2d$, as illustrated in Fig. 14. However, if satellites are deployed at relatively short mutual distances then their coverage areas will be reduced on the ground so that their antennas will point more directly downwards leading to greater discrimination between the satellite to ground directions and the satellite to satellite directions.

For operation of a satellite based mobile communication system, it will clearly be necessary to set up calls. This requires a broadcast channel (BCH) and a random access channel (RACH). These can be provided by allocating a downlink time slot to the BCH and the immediately preceding uplink time slot to the RACH channel. Data is repeated on the broadcast channel so it is not essential for every terminal to receive every BCH time slot. However, a terminal could be assigned an uplink time slot that requires it to transmit during the reception period of the BCH time slot. For this reason it is proposed to allocate *two* BCH time slots separated by $\frac{1}{2}$ of a TDMA repeat period. Then *every* terminal will be able to receive one or other BCH time slot.